Grid, Cluster and Cloud Computing

K. Keahey[1], D. Laforenza[1], A. Reinefeld[2], P. Ritrovato[2], D. Thain[2], and N. Wilkins-Diehr[2]

[1] Topic Chairs
[2] Members

Grid computing is a major research area with strong involvement from both academia and the computing industry. The common vision is that grid computing represents the culmination of truly general distributed computing across various resources in a ubiquitous, open-ended infrastructure to support a wide range of different application areas. Although significant progress has been made in the design and deployment of grids, many challenges still remain before the goal of a user-friendly, efficient, and reliable grid can be realized. Grid research issues cover many areas of computer science to address the fundamental capabilities and services that are required in a heterogeneous environment, such as adaptability, scalability, reliability and security, and to support applications as diverse as ubiquitous local services, enterprise-scale virtual organizations, and internet-scale distributed supercomputing. Cloud computing is also emerging as an alternate platform for large-scale distributed applications where resources are typically provided by a single administrative domain in a pay per-use mode. To some, cloud computing is a natural evolution of grid computing, to others, it is a complementary and perhaps competing technology. Grid and cloud research will greatly benefit from interactions with the many related areas of computer science, making Euro-Par an excellent venue to present results and discuss issues.

This year, 29 papers discussing some those issues were submitted to this topic. Each paper was reviewed by at least three reviewers and, finally, we were able to select 8 regular papers (28% acceptance rate). The accepted papers discuss interesting challenges arising in grid and cloud computing.

In particular, in the paper "What Is the Price of Simplicity? A Cross-platform Evaluation of the SAGA API" M. den Burger et al. describe the effects on expressiveness and ease of programming using abstractions defined by the SAGA API (Simple API for Grid Applications) and shows what price is paid for the abstraction in terms of performance.

Two papers focus on examining the distribution challenge in grids and clouds. "Using Network Information to Perform Meta-scheduling in Advance in Grids" by L. Toms et al. discusses the challenges of providing quality of service (QoS) in the grid environment and proposes a system that provides it through network-aware job scheduling in advance while "Deployment of a hierarchical middleware" by E. Caron et al. proposes a hierarchical model for gird middleware and an algorithm for automatically and efficiently deploying it is a distributed environment taking into account throughput ratios.

P. D'Ambra, M. Guarracino, and D. Talia (Eds.): Euro-Par 2010, Part I, LNCS 6271, pp. 341–342, 2010.

Investigating the properties of workflow systems continues to be a significant challenge in the grids. In "Toward Real-time, Many-Task Applications on Large Distributed Systems", S. Yi et al. design and implement a real-time task management system for many-task computing - the system builds on the BOINC platform but provides impressive O(1) worst-case execution time for task management operations. "Scheduling Scientific Workflows to Meet Soft Deadlines in the Absence of Failure Models" by K. Plankensteiner et al., proposes a dynamic execution and scheduling heuristic able to schedule workflow applications with a high degree of fault tolerance, while taking into account soft deadlines. And finally "User-centric, Heuristic Optimization of Service Composition in Clouds" by K. Kofler at al. proposes a system to map costumer requirements onto functional and non-functional attributes of the services and evaluates the resulting system in the context of the Kepler workflow tool.

Challenges arising in cloud computing have been investigated all throughout the programming stack. The paper "A GPGPU transparent virtualization component for high performance computing clouds" by G. Giunta et al. focuses on the minutiae of virtualization and explains how a virtual machine can access GPGPUs in a transparent way, with an overhead only slightly greater than a real machine/GPGPU setup whereas the paper "A Distributed Market Framework for Large-scale Resource Sharing" by M. Mihailescu et al. addresses things on a higher level and presents scalable distributed market framework for the allocation of shared resources in large distributed systems.

We would like to take the opportunity of thanking all authors who shared their results with EuroPar community as well as the Euro-Par Organizing Committee, and the referees for their hard work and insightful comments, whose efforts have made this conference and this topic possible.